

CLAIMS

1. A method of processing data in a computer system comprising at least one host and at least one storage system which stores data for the at least one host, the method comprising acts of:
 - (a) receiving, at the storage system, a request from the at least one host to access a unit of data stored on the at least one storage system; and
 - (b) in response to the request, returning to the at least one host information related to where the unit of data is physically stored on the at least one storage system.
2. The method of claim 1, wherein the storage system is a content addressable storage system and the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, and wherein the request received in the act (a) identifies the unit of data using a content address.
3. The method of claim 1, further comprising an act (c), subsequent to the act (b) of receiving, at the at least one storage system, a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the storage system.
4. The method of claim 1, wherein the act (a) further comprises an act of receiving a query request that specifies at least one criterion and requests units of data stored on the storage system that satisfy the at least one criterion based.
5. The method of claim 4, wherein the at least one criterion includes a time period in which units of data were stored on the storage system.
6. The method of claim 1, wherein the request to access the unit of data includes a request to write the unit of data, and wherein the act (a) further comprises an act of receiving the unit of data from the host.

7. The method of claim 6, further comprising an act of returning, to the host, an address associated with the unit of data.
8. The method of claim 7, further comprising an act (c), subsequent to the act (b), of receiving, at the at least one storage system, a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the storage system..
9. The method of claim 1, wherein the storage system includes a plurality of storage nodes, and wherein the information related to where the unit of data is physically stored on the storage system includes information identifying the ones of the plurality of storage nodes on which the unit of data is stored.
10. The method of claim 9, wherein at least one of the plurality of storage nodes includes a plurality of disks, and wherein the information related to where the unit of data is physically stored on the storage system includes information related to on which of the plurality of disks the unit of data is stored.
11. At least one computer readable medium encoded with instructions that, when executed on a computer system perform a method of processing data, wherein the computer system comprises at least one host and at least one storage system which stores data for the at least one host, the method comprising acts of:
 - (a) receiving, at the storage system, a request from the at least one host to access a unit of data stored on the at least one storage system; and
 - (b) in response to the request, returning to the at least one host information related to where the unit of data is physically stored on the at least one storage system.
12. The at least one computer readable medium of claim 11, wherein the storage system is a content addressable storage system and the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, and wherein the request received in the act (a) identifies the unit of data using a content address.

13. The at least one computer readable medium of claim 11, wherein the method further comprises an act (c), subsequent to the act (b) of receiving, at the at least one storage system, a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the storage system.
14. The at least one computer readable medium of claim 13, wherein the act (a) further comprises an act of receiving a query request that specifies at least one criterion and requests units of data stored on the storage system that satisfy the at least one criterion.
15. The at least one computer readable medium of claim 14, wherein the at least one criterion includes a time period in which units of data were stored on the storage system.
16. The at least one computer readable medium of claim 11, wherein the request to access the unit of data includes a request to write the unit of data, and wherein the act (a) further comprises an act of receiving the unit of data from the host.
17. The at least one computer readable medium of claim of claim 11, wherein the method further comprises an act of returning, to the host, an address associated with the unit of data.
18. The at least one computer readable medium of claim 17, wherein the method further comprises an act (c), subsequent to the act (b), of receiving, at the at least one storage system, a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the storage system.
19. The at least one computer readable medium of claim 11, wherein the storage system includes a plurality of storage nodes, and wherein the information related to

where the unit of data is physically stored on the storage system includes information identifying the ones of the plurality of storage nodes on which the unit of data is stored.

20. The at least one computer readable medium of claim 19, wherein at least one of the plurality of storage nodes includes a plurality of disks, and wherein the information related to where the unit of data is physically stored on the storage system includes information related to on which of the plurality of disks the unit of data is stored.

21. A storage system for use in a computer system that includes at least one host, wherein the storage system stores data for the at least one host, the storage system comprising:

at least one storage device to store data received from the at least one host; and

at least one controller that receives a request from the at least one host to access a unit of data stored on the storage system and in response to the request, returns to the at least one host information related to where the unit of data is physically stored on the storage system.

22. The storage system of claim 21, wherein the storage system is a content addressable storage system, wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, and wherein the request from the at least one host identifies the unit of data using a content address.

23. The storage system of claim 21, wherein the at least one controller, subsequent to returning to the at least one host information related to where the unit of data is physically stored, receives a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the content addressable storage system.

24. The storage system of claim 23, wherein the at least one controller receives a query request that specifies at least one criterion and requests units of data stored on the

storage system that satisfy the at least one criterion, and wherein the at least one controller processes the request.

25. The storage system of claim 24, wherein the at least one criterion includes a time period in which units of data were stored on the storage system.

26. The storage system of claim 21, wherein the request to access the unit of data includes a request to write the unit of data, and wherein the at least one controller receives the unit of data from the at least one host.

27. The storage system of claim 26, wherein the at least one controller, subsequent to returning to the at least one host information related to where the unit of data is physically stored, returns to the at least one host an address associated with the unit of data.

28. The storage system of claim 27, wherein the at least one controller, receives a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the storage system.

29. The storage system of claim 21, further comprising a plurality of storage nodes, wherein the information related to where the unit of data is physically stored on the content addressable storage system includes information identifying the ones of the plurality of storage nodes on which the unit of data is stored.

30. The storage system of claim 29, wherein at least one of the plurality of storage nodes includes a plurality of disks, and wherein the information related to where the unit of data is physically stored on the content addressable storage system includes information related to on which of the plurality of disks the unit of data is stored.

31. A method of processing data in a computer system comprising at least one host and at least one storage system which stores data for the at least one host, the method comprising acts of:

(a) sending, to the at least one storage system, a request to access a unit of data stored on the storage system; and

(b) receiving, from the at least one storage system, information related to where the unit of data is physically stored on the at least one storage system.

32. The method of claim 31, wherein the at least one storage system is a content addressable storage system and the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, and wherein the request sent in the act (a) identifies the unit of data using a content address.

33. The method of claim 32, further comprising an act (c), subsequent to the act (b), of sending to the at least one storage system a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the at least one storage system.

34. The method of claim 33, wherein the act (a) further comprises an act of sending, to the at least one storage system, a query request that specifies at least one criterion and requests units of data stored on the storage system that satisfy the at least one criterion based on the metadata.

35. The method of claim 34, wherein the at least one criterion includes a time period in which units of data were stored on the storage system.

36. The method of claim 31, wherein the request to access the unit of data includes a request to write the unit of data, and wherein the act (a) further comprises an act of sending the unit of data to the at least one storage system.

37. The method of claim 33, wherein the request to access the unit of data includes a request to read the unit of data, and wherein the method further comprises an act of receiving the unit of data from the at least one storage system.

38. The method of claim 37, wherein the act (a) further comprises an act of sending, to the at least one storage system, an address associated with the unit of data.

39. At least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method of processing data, wherein the computer system comprises at least one host and at least one storage system which stores data for the at least one host, the method comprising acts of:

(a) sending, to the at least one storage system, a request to access a unit of data stored on the storage system; and

(b) receiving, from the at least one storage system, information related to where the unit of data is physically stored on the at least one storage system.

40. The at least one computer readable medium of claim 39, wherein the at least one storage system is a content addressable storage system and the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, and wherein the request sent in the act (a) identifies the unit of data using a content address.

41. The at least one computer readable medium of claim 39, wherein the method further comprises an act (c), subsequent to the act (b), of sending, to the at least one storage system a second request to access the unit of data, the second request including the information related to where the unit of data is physically stored on the at least one storage system.

42. The at least one computer readable medium of claim 41, wherein the act (a) further comprises an act of sending, to the at least one storage system, a query request that specifies at least one criterion and requests units of data stored on the storage system that satisfy the at least one criterion.

43. The at least one computer readable medium of claim 42, wherein the at least one criterion includes a time period in which units of data were stored on the storage system.

44. The at least one computer readable medium of claim 39, wherein the request to access the unit of data includes a request to write the unit of data, and wherein the act (a) further comprises an act of sending the unit of data to the at least one storage system.

45. The at least one computer readable medium of claim 41, wherein the request to access the unit of data includes a request to read the unit of data, and wherein the method further comprises an act of receiving the unit of data from the at least one storage system.

46. The at least one computer readable medium of claim 45, wherein the act (a) further comprises an act of sending, to the at least one storage system, an address associated with the unit of data.

47. A host computer for use in a computer system that includes the host computer and at least one storage system, the host computer comprising:
 at least one storage device; and
 at least one controller that:
 sends, to the at least one storage system, a request to access a unit of data stored on the at least one storage system; and
 in response to the request, receives from the at least one storage system, information related to where the unit of data is physically stored on the at least one storage system and stores the information in the at least one storage device.

48. The host computer of claim 47, wherein the at least one storage system is a content addressable storage system and the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, and wherein the request sent to the at least one storage system identifies the unit of data using a content address.

49. The host computer of claim 47, wherein the at least one controller, subsequent to receiving the information related to where the unit of data is physically stored, sends, to the at least one storage system, a second request to access the unit of data, the second

request including the information related to where the unit of data is physically stored on the at least one storage system.

50. The host computer of claim 49, wherein the request to access the unit of data includes a query request that specifies at least one criterion and requests units of data stored on the storage system that satisfy the at least one criterion.

51. The host computer of claim 50, wherein the at least one criterion includes a time period in which units of data were stored on the storage system.

52. The host computer of claim 47, wherein the request to access the unit of data includes a request to write the unit of data, and wherein the controller sends the unit of data to the at least one storage system.

53. The host computer of claim 47, wherein the request to access the unit of data includes a request to read the unit of data, and wherein the controller sends the unit of data to the at least one storage system.

54. The host computer of claim 53, wherein the controller, sends to the at least one storage system an address associated with the unit of data.

55. A method of processing data in a computer system comprising at least one host and at least one storage system, the method comprising acts of:

(a) receiving, at the at least one storage system, a request from the at least one host to access a unit of data, the request having an identifier that can be used to access the unit of data; and

(b) in response to the request, returning to the at least one host an identifier that can be used to access the unit of data and additional information, separate from the identifier, related to where the unit of data is physically stored on the at least one storage system.

56. The method of claim 55, wherein the request to access the unit of data includes a request to write the unit of data and the act (a) further comprises acts of receiving the unit of data from the host and storing the unit of data at a storage location of the storage system.

57. The method of claim 55, further comprising an act of receiving a second request to access the unit of data, the second request including the identifier associated with the unit of data and the information related to where the data is physically stored on the storage system.

58. The method of claim 55, wherein the at least one storage system includes a plurality of storage nodes and wherein the information related to where the unit of data is physically stored on the at least one storage system includes information related to on which of the plurality of storage nodes the unit of data is stored.

59. The method of claim 55, wherein the at least one storage system includes a plurality of storage nodes, at least some of the plurality of storage nodes including a plurality of disks, and wherein the information related to where the unit of data is physically stored on the storage system includes information related to on which of the plurality of disks the unit of data is stored.

60. At least one computer readable medium encoded with instructions that, when executed on a computer system perform a method of processing data, wherein the computer system comprises at least one host and at least one storage system, the method comprising acts of:

(a) receiving, at the at least one storage system, a request from the at least one host to access a unit of data, the request having an identifier that can be used to access the unit of data; and

(b) in response to the request, returning to the at least one host an identifier that can be used to access the unit of data and additional information, separate from the identifier, related to where the unit of data is physically stored on the at least one storage system.

61. The at least one computer readable medium of claim 60, wherein the request to access the unit of data includes a request to write the unit of data and the act (a) further comprises acts of receiving the unit of data from the host and storing the unit of data at a storage location of the storage system.

62. The at least one computer readable medium of claim 60, wherein the method further comprises an act of receiving a second request to access the unit of data, the second request including the identifier associated with the unit of data and the information related to where the data is physically stored on the storage system.

63. The at least one computer readable medium of claim 60, wherein the at least one storage system includes a plurality of storage nodes and wherein the information related to where the unit of data is physically stored on the at least one storage system includes information related to on which of the plurality of storage nodes the unit of data is stored.

64. The at least one computer readable medium of claim 60, wherein the at least one storage system includes a plurality of storage nodes, at least some of the plurality of storage nodes including a plurality of disks, and wherein the information related to where the unit of data is physically stored on the storage system includes information related to on which of the plurality of disks the unit of data is stored.

65. A storage system for use in a computer system, including the storage system and at least one host, the storage system comprising:

at least one storage device to store data received from the at least one host; and

at least one controller that:

receives, at the at least one storage system, a request from the at least one host to access a unit of data, the request having an identifier that can be used to access the unit of data; and

in response to the request, returns to the at least one host an identifier that can be used to access the unit of data and additional information, separate from the

identifier, related to where the unit of data is physically stored on the at least one storage system.

66. The storage system of claim 65, wherein the request to access the unit of data includes a request to write the unit of data, and the controller receives the unit of data from the host and stores the unit of data on the at least one storage device.

67. The storage system of claim 65, wherein the controller receives a second request to access the unit of data, the second request including the identifier associated with the unit of data and the information related to where the data is physically stored on the storage system.

68. The storage system of claim 65, wherein the at least one storage device comprises a plurality of storage nodes, wherein the information related to where the unit of data is physically stored on the at least one storage system includes information related to on which of the plurality of storage nodes the unit of data is stored.

69. The storage system of claim 65, wherein the at least one storage device comprises a plurality of storage nodes, at least some of the plurality of storage nodes including a plurality of disks, wherein the information related to where the unit of data is physically stored on the storage system includes information related to on which of the plurality of disks the unit of data is stored.

70. A method of processing data in a computer system comprising at least one host and at least one storage system, the method comprising acts of:

(a) sending, to the at least one storage system, a request from the at least one host to access a unit of data, the request having an identifier that can be used to access the unit of data; and

(b) receiving, from the at least one storage system, an identifier that can be used to access the unit of data and additional information, separate from the identifier, related to where the unit of data is physically stored on the at least one storage system.

71. The method of claim 70, wherein the request to access the unit of data includes a request to write the unit of data and the act (a) further comprises acts of sending the unit of data from the at least one host for storage at a storage location of the storage system.

72. The method of claim 70, further comprising an act of sending a second request to access the unit of data from the at least one host to the at least one storage system, the second request including the identifier associated with the unit of data and the information related to where the data is physically stored on the storage system.

73. At least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method of processing data, wherein the computer system comprises at least one host and at least one storage system, the method comprising acts of:

(a) sending, to the at least one storage system, a request from the at least one host to access a unit of data, the request having an identifier that can be used to access the unit of data; and

(b) receiving, from the at least one storage system, an identifier that can be used to access the unit of data and additional information, separate from the identifier, related to where the unit of data is physically stored on the at least one storage system.

74. The at least one computer readable medium of claim 73, wherein the request to access the unit of data includes a request to write the unit of data and the act (a) further comprises acts of sending the unit of data from the at least one host for storage at a storage location of the storage system.

75. The at least one computer readable medium of claim 73, wherein the method further comprises an act of sending a second request to access the unit of data from the at least one host to the at least one storage system, the second request including the identifier associated with the unit of data and the information related to where the data is physically stored on the storage system.

76. A host computer for use in a computer system including the host computer and at least one storage system, the host computer comprising:

at least one storage device; and

at least one controller that:

sends, to the at least one storage system, a request from the at least one host to access a unit of data, the request having an identifier that can be used to access the unit of data; and

in response to the request, receives from the at least one storage system an identifier that can be used to access the unit of data and additional information, separate from the identifier, related to where the unit of data is physically stored on the at least one storage system, and stores the additional information on the at least one storage device.

77. The host computer of claim 76, wherein the request to access the unit of data includes a request to write the unit of data and the at least one controller sends the unit of data from the at least one host to the at least one storage system.

78. The host computer of claim 76, wherein the at least one controller is adapted to send, along with a second request to access the unit of data from the at least one host to the at least one storage system, the identifier associated with the unit of data and the information related to where the data is physically stored on the storage system.